AD-A047 028 FLORIDA UNIV GAINESVILLE DEPT OF INDUSTRIAL AND SYS--ETC F/G 12/1
AN EFFICIENT COMPUTATIONAL ALTERNATIVE TO 'USING LINEAR PROGRAM--ETC(U)
OCT 77 L E DANIEL, S HART, T J HODGSON N00014-76-C-0096

UNCLASSIFIED RR-77-9 NL

1 OF 1
ADA047028

END
DATE
FILMED
1 -78
DDC

RESEARCH REPORT

Industrial & Systems
Engineering Department
University of Florida
Gainesville, FL. 32611

AN EFFICIENT COMPUTATIONAL ALTERNATIVE TO
'USING LINEAR PROGRAMMING TO DESIGN
OIL POLLUTION DETECTION SCHEDULES'

Research Report No. 77-9

by

Lee E. Daniel, Jr.*
Sandal Hart**
Thom J. Hodgson+

October, 1977

+Department of Industrial & Systems Engineering
University of Florida
Gainesville, Florida 32611

This research was supported in part by the Office of Naval Research, under contract number N00014-76-C-0096.

*Vector Research, Inc., Ann Arbor, Michigan 48106.

**Management Systems Division, Proctor and Gamble, Inc.
Cincinnati, Ohio 45222

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

14

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER RR-77-9	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) 6 An Efficient Computational Alternative to 'Using Linear Programming to Design Oil Pollution Detection Schedules',		5. TYPE OF REPORT & PERIOD COVERED 9 Research rept. 6. PERFORMING ORG. REPORT NUMBER 77-9
7. AUTHOR(s) 10 Lee E. Daniel, Jr., Sandal Hart Thom J. Hodgson		8. CONTRACT OR GRANT NUMBER(s) 15 NØ0014-76-C-Ø096
9. PERFORMING ORGANIZATION NAME AND ADDRESS Industrial & Systems Engineering University of Florida Gainesville, FL 32611		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Office of Naval Research Arlington, VA		12. REPORT DATE 11 October, 1977 13. NUMBER OF PAGES 15
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) 12 16p.		15. SECURITY CLASS. (of this report) Unclassified 15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Bloc[illegible]ferent from Report)

N/A

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Markov Decision Processes
Semi-Markov Processes
Pollution Detection
Search
Dynamic Programming

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

In Olson, Wright, and McKell's recent paper on the design of oil pollution detection schedules, an interesting and inventive development and application of a Markov Decision Process was presented. Optimal schedules for patrol flights of surveillance aircraft were found using linear programming. In this paper the model has been reformulated as a discrete time semi-Markov process. Significant computational advantages accrue from this alternative approach.

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

404399

Abstract

An Efficient Computational Alternative To
'Using Linear Programming to Design
Oil Pollution Detection Schedules'

by

Lee E. Daniel, Jr.
Sandal S. Hart
Thom J. Hodgson

In Olson, Wright, and McKell's recent paper on the design of oil pollution detection schedules, an interesting and inventive development and application of a Markov Decision Process was presented. Optimal schedules for patrol flights of surveillance aircraft were found using linear programming. In this paper the model has been reformulated as a discrete time semi-Markov process. Significant computational advantages accrue from this alternative approach.

Olson, Wright, and McKell [6], recently reported on a very interesting application of the linear programming formulation for Markov decision processes. The U. S. Coast Guard is charged with the responsibility for environmental protection in our coastal areas. Specifically, they are chartered to operate electronic sensor equipped patrol aircraft whose mission, in this case, is the detection and prevention of oil and hazardous material pollution in coastal and offshore areas. A search model was developed for scheduling patrol flights. This model is part of a system called Pollution Detection and Prevention System (PDAPS). The objective of the model is to find a flight plan or schedule which maximizes the expected number of pollution detections per patrol flight. Each flight may cover a given number of known geographical sectors where pollution is likely to occur. The probability of a pollution incident occurring in a geographical area is obtained from historical pollution statistics, shipping statistics, and pollution prediction models which are contained in PDAPS. It may not be possible to search all sectors of interest in one flight. Of those sectors which are searched, there may be multiple possible flight patterns depending on physical properties of the sector and flight altitude. For example, three flight patterns are shown for a sector in figure 1. The detection probability for a sector varies according to the pattern flown.

In addition to maximizing the expected number of pollution incidents detected, the pollution flight schedules include a randomness factor in order to have a preventive effect on intentional polluters. By this we mean that schedule generation is performed in two stages. First, an optimal "expected value"schedule is generated. Second, each time an actual flight is made, a schedule is generated randomly from the "expected value" schedule. The amount of "randomness" in the actual schedule is related to a randomness factor ε, to be defined shortly.

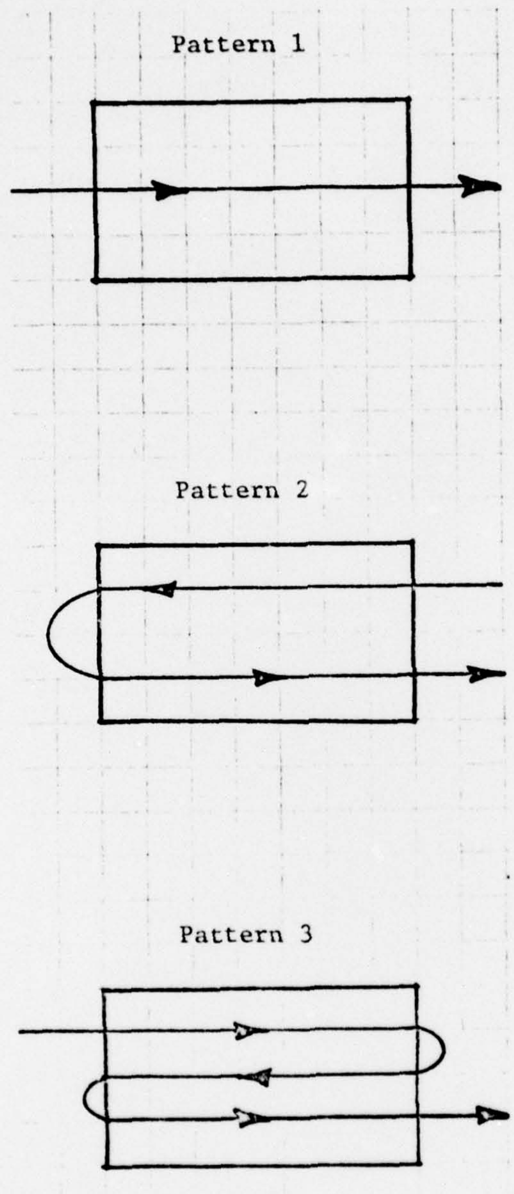

Figure 2

Three Different Flight Patterns
for a Geographical Pollution Sector

States of the system, for modelling purposes, are described by a 3-tuple in which the first entry designates the geographical sector, the second indicates the sector exit point (i.e. the flight pattern used to search the sector); and the third gives the remaining flight time. It is clear that the state definition has the Markov property, in that for any given system state (i.e., the location of the aircraft and the remaining flight time) the definition includes the information necessary to plan the rest of a flight. The reader may note that the definition, however, does allow revisiting of geographic sectors (i.e., subtours). This shortcoming is easily eliminated in practice. The reader is referred to [5, 6] for more detail. Included in the set of states are states representing both the beginning and the end of the flight (normally the same location, but not necessarily).

We now move to a linear programming formulation of the patrol flight scheduling problem. The objective of the linear program is, for a given randomness factor ε, to maximize the expected number of pollution detections. Note that in the following discussion the 3-tuple system state variable is represented by a single dimensioned variable. Define S_i as the set of possible successor states associated with state i. If S denotes the set of all states then $S = S_1 \cup S_2 \cup \cdots \cup S_N$ where N is the number of states. Let M denote the total flight time available for a patrol mission. τ_{ij} is the total flight time associated with transition from state i to state j. Pd_{ij} is the probability of detecting a pollution incident associated with the transition from state i to state j. The introduction of randomness into the flight schedules is accomplished in the following manner. Let $q_{ij}(a)$ be the probability of going to state j given that the current state is i and it is desired to transition to state a. This is related to the randomness factor ε as follows:

$0 \leq \varepsilon \leq 1$, $N(S_i)$ = cardinality of S_i

$$q_{ij}(a) = \begin{cases} 0 & , j \notin S_i \\ 1 - \varepsilon & , j = a \\ \dfrac{\varepsilon}{N(S_i) - 1} & , j \in S_i,\ j \neq a. \end{cases}$$

Note that

$$c_i^a = \sum_{j \in S_i} q_{ij}(a) Pd_{ij}$$

represents the expected number of detections from state i given that the decision is to go to a. Let z_{ia} be the probability of being in state i and choosing to go to state a.

The final linear programming model (equivalent to model II in [6]) is: Find $\{z_{ia}\}$ $i \in S$, $a \in S_i$ in order to

$$\max \sum_{i \in S} \sum_{a \in S_i} c_i^a z_{ia},$$

Subject to

$$\sum_{i \in S} \sum_{a \in S_i} z_{ia} = 1$$

$$\sum_{i \in S} \sum_{a \in S_i} z_{ia}(\delta_{ij} - q_{ij}(a)) = 0, \quad \text{all } j \in S,$$

where

$$\delta_{ij} = \begin{cases} 0 & , i \neq j \\ 1 & , i = j \end{cases}$$

It should be noted that the linear programming formulation of Olson, et al. [6], is equivalent to that of Derman [1], Wolfe and Dantzig [8], and Manne [4] for

Markov decision processes. The structure of the linear program insures that one and only one policy is chosen for each state of the system. The execution time for solving this system as a linear program is quite large, particularly for problems of practical size (see Table 1).

In reviewing the process, it can be observed that, disregarding the remaining flight time, the process of going from sector to sector is Markovian. Therefore, the process can be viewed as a discrete time semi-Markov process where the state of the system is described by a 2-tuple in which the first entry is the geographic sector and the second indicates the sector exit point (flight pattern). The time between transitions is the transit time between sector exit points. The semi-Markov structure results in a considerable reduction in the number of states of the system. Since the problem has a finite planning horizon (length of the flight), dynamic programming appears to be the appropriate solution technique. From Howard [3], the general form of the finite horizon, discrete time semi-Markov dynamic recursion is (ignoring boundary conditions):

$$v_i(n) = \max_a \{c_i^a + \sum_{j=1}^{N} q_{ij}(a) \sum_{m=1}^{n} v_j(n - m)\, h_{ij}^a(m)\}, \tag{1}$$

where

$v_i(n)$ is the maximum expected number of pollution detections over the remaining n time units given the process starts in state i.

$h_{ij}^a(m)$ is the probability that m time units will be required to go from state i to state j when it is desired to transition to state a.

Note that the state designations i and j now refer to the redefined 2-dimensional states of the semi-Markov model. For this particular problem, the inner sum over the transition time probabilities in (1) has only one term since the transition (flight) time between states is assumed to be deterministic. Hence, there is a matrix of transition times from state i to state j (τ_{ij}), rather than one of functions of transition times from state i to state j ($h_{ij}(m)$). Equation (1) simplifies to the following form:

Table 1. Computation Times for Linear Program***

Length of flight (minutes)	Number of Markov States	Computation Time* (Seconds)
**	66	21.
120	341	552.
420	1565	7815.

[* CPU Seconds on CDC 6000 Series Computer [5]]
[** Not able to determine]
[*** Linear Programs run using the CDC OPTIMA System]

$$v_i(n) = \max_a \{c_i^a + \sum_{j=1}^{N} q_{ij}(a) v_j(n - \tau_{ij})\}, \tag{2}$$

where τ_{ij} is the transition time from state i to state j. The procedure is initiated by setting

$v_i(0) = -\infty$, $\quad$ $i \neq$ the final (or home) state,

$v_i(0) = 0$, $\quad$ $i =$ the final (or home) state, and

$v_i(n) = -\infty$, $\quad$ for negative values of n

The procedure is terminated when $v_I(M)$ is calculated, where I is the initial (or beginning) state and M is the length of the flight.

It should be noted that the semi-Markov formulation is equivalent to that of Olson, et al., and the state reduction techniques discussed in [5, 6] apply to this formulation. Computational advantages accrue from the significant reduction in the number of states in the system and the finite horizon dynamic programming approach (as opposed to the infinite horizon LP approach).

Equivalence of size, for comparison of computational efficiency of the semi-Markov and Markov formulations, is easily established.
Certain entries of the $v_i(n)$ matrix can be determined to be infeasible (i.e., it is either impossible to reach state i from the start point in n time units, or it is impossible to reach the finish point from state i in the flight time remaining). In addition, it is possible to limit, artificially, the time window within which each sector can be visited. For instance, due to operational considerations associated with a given data set to be used as input for the dynamic program, it may be apparent that a certain geographic sector can only be visited early in the flight, if at all. It would make sense, then, to declare those entries of $v_i(n)$ associated with the geographic sector to be infeasible for time periods (n) greater than the latest reasonable visitation time. The remaining (feasible) elements of $v_i(n)$ each represent states of the system for the Markov formulation. Therefore, counting the feasible entries of the $v_i(n)$ matrix gives the equivalent number of states of the Markov problem for a given semi-Markov problem.

The data used by Olson et al., was not available. Therefore, a problem was formulated using realistic data from [7]. The sectors considered were potential oil well drilling sites and shipping lanes in the Gulf of Mexico off the Florida coast. The aircraft was assumed to fly at a speed of 130 knots. Flight patterns were designed for each sector and the detection probabilities were randomly generated. There were 35 sectors which, when combined with the various flight patterns, resulted in 111 semi-Markov states. Time was discretized, as in [5, 6], in minutes. The equivalent number of Markov states depended on the length of the flight and the extent to which various state reduction techniques [5, 6] were applied. The computational results for the various runs made are displayed in Table 2. In addition, the computational results are included in Figure 2 for comparison with the linear programming results. (Note that figure 2 is a semi-log graph).

The relative difference in computation times is quite large. It might be argued that some of the difference can be attributed to the relative speeds of the CDC 6000 series computer used for the Linear Program and the IBM-370 model 165 computer used for the Dynamic Program, but the computational differences (more than 4 orders of magnitude) are large enough to absorb easily any differences in machine speed. For the problem at hand, the power of the discrete time semi-Markov process as a modelling tool is that it brings forth the underlying structure in a more straightforward fashion. This allows a simplified computational approach to the optimization and results in the computational efficiencies observed.

The computer storage requirements for the experimental dynamic program code are reasonably modest, all things considered. The 6 hr. (360 minute) flight problem (the largest we ran) requires less than 260K bytes of core. No additional off-line storage is necessary. Rather large efficiencies (~50%) could be affected through list processing to eliminate storage requirements for infeasible entries of $v_i(n)$, but that was not implemented in the code.

Table 2. Computation Times for Dynamic Program**

Number of Semi-Markov States	Length of flight (minutes)	Equivalent Number of Markov States	Computation Time* (Seconds)
111	180	1076	.396
111	180	2914	1.151
111	360	4120	1.694
111	360	7303	2.626
111	360	14592	7.125

[*CPU Seconds on IBM 370-165 computer]
[** Dynamic Program coded in FORTRAN IV]

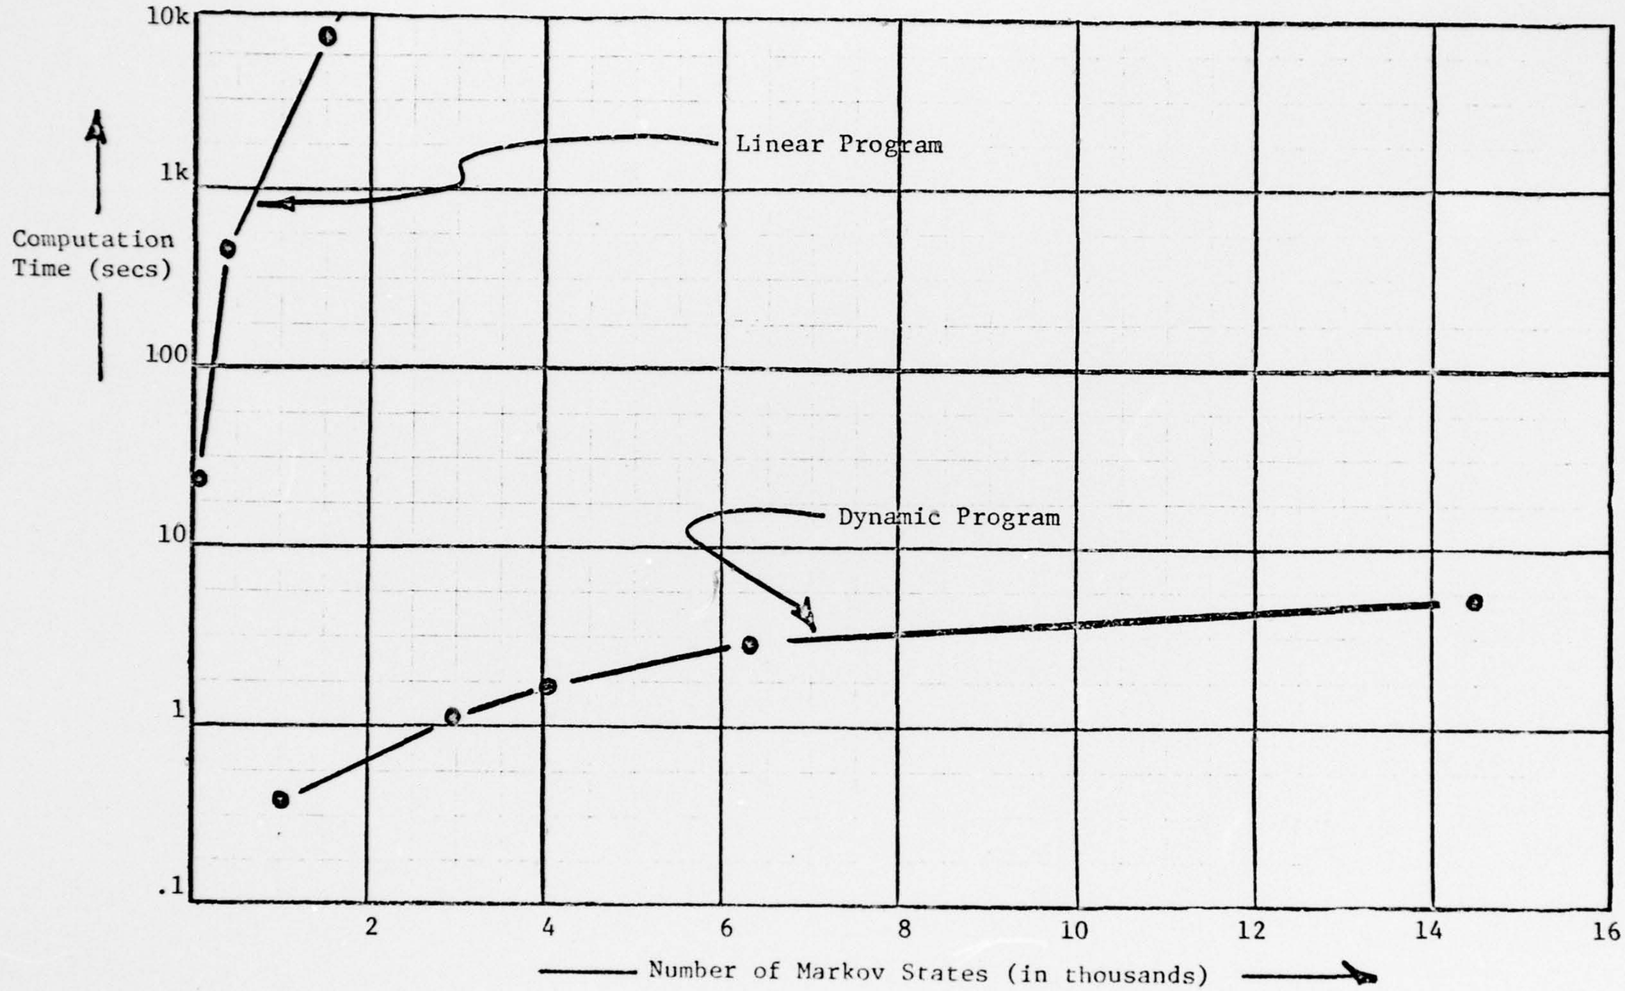

Figure 1. Computation Times for Linear Program and Dynamic Program.

References

1. Derman, C., "On Sequential Decisions and Markov Chains," Management Science, Vol. 9, pp. 16-24, 1962

2. Howard, Ronald A., Dynamic Programming and Markov Processes, The Massachusetts Institute of Technology Press, Cambridge, Mass., 1960.

3. Howard, Ronald A., Dynamic Probabilistic Systems Volume II: Semi-Markov and Decision Processes, John Wiley and Sons, Inc., New York, 1971.

4. Manne, Alan S., "Linear Programming and Sequential Decisions," Management Science, Vol. 6, pp. 259-269, 1960.

5. McKell, Lynn J., "A Computer Based Pollution Surveillance Scheduling System," unpublished Ph.D. dissertation, December, 1973, Purdue University.

6. Olson, David G., Wright, Gordon P., and McKell, Lynn J., "Using Linear Programming to Design Oil Pollution Detection Schedules," AIIE Transactions, Vol. 8, No. 3, September, 1976.

7. Welsh, Leonard W., Jr., "Oil Spill Simulation in the Florida Gulf Outer Continential Shelf," unpublished Master's Thesis, University of Florida, 1976.

8. Wolfe, Philip and Dantzig, G. B., "Linear Programming in a Markov Chain," Operations Research, Vol. 10, No. 5, pp. 702-710, September-October, 1962.